

WHAT IS CLAIMED IS:

1. A method for adding an acoustic description of a word to a speech recognition lexicon, the method comprising:

converting the text of the word into at least one orthographically derived acoustic description of the word;

generating a score for an orthographically derived acoustic description based in part on a comparison between the orthographically derived acoustic description and a speech signal representing a user's pronunciation of the word;

decoding the speech signal representing the user's pronunciation of the word to produce a decoded acoustic description of the word and a score for the decoded acoustic description; and

selecting one of the orthographically derived acoustic description and the decoded acoustic description as the acoustic description of the word based on the score for the orthographically derived acoustic description and the score for the decoded acoustic description.

2. The method of claim 1 wherein generating a score for an orthographically derived acoustic description comprises generating an acoustic model score.

3. The method of claim 2 wherein decoding the speech signal comprises generating an acoustic model score for at least one decoded acoustic description and using the score as at least part of the score for the decoded acoustic description.

4. The method of claim 3 wherein generating an acoustic model score for the orthographically derived acoustic description and generating an acoustic model score for at least one decoded acoustic description comprises using the same acoustic model to generate both acoustic model scores.

5. The method of claim 3 wherein decoding the speech signal further comprises generating a language model score for the at least one decoded acoustic description and using the language model score as part of the score for the at least one decoded acoustic description.

6. The method of claim 5 wherein generating an acoustic model score and generating a language model score for at least one decoded acoustic description comprises generating an acoustic model score and a language model score for a sequence of syllable-like units and wherein the decoded acoustic description is derived from the sequence of syllable-like units.

7. The method of claim 6 wherein deriving the decoded acoustic description from the sequence of syllable-like units comprises dividing the sequence of syllable-like units into a sequence of phonemes.

8. The method of claim 6 wherein generating a language model score comprises generating a language model score based on a trigram language model for syllable-like units.

9. The method of claim 6 wherein generating an acoustic model score for a sequence of syllable-like units comprises generating acoustic model scores for each of a sequence of phonemes that form the sequence of syllable-like units.

10. The method of claim 1 further comprising displaying a user interface comprising an edit box in which a user may enter the text of the word and a list box that displays words for which an acoustic description has been previously added to the speech recognition lexicon.

11. The method of claim 10 further comprising:
receiving an indication that a user has selected a word in the list box;
retrieving the added acoustic description of the word from the speech recognition lexicon; and
converting the retrieved acoustic description into an audible signal.

12. A computer-readable medium having computer-executable instructions for performing steps comprising:

receiving text of a word for which a phonetic description is to be added to a speech recognition lexicon;

receiving a representation of a speech signal produced by a person pronouncing the word;

converting the text of the word into a text-based phonetic description of the word;

generating a speech-based phonetic description of the word from the representation of the speech signal; and

selecting a phonetic description of the word to add to the speech recognition lexicon by selecting between the text-based phonetic description and the speech-based phonetic description based in part on the correspondence between each phonetic description and the representation of the speech signal.

13. The computer-readable medium of claim 12 wherein generating a speech-based phonetic description comprises:

generating a plurality of possible phonetic descriptions;

using at least one model to score each possible phonetic description; and

selecting the possible phonetic description with the highest score as the speech-based phonetic description.

14. The computer-readable medium of claim 13 wherein using at least one model comprises using an acoustic model and a language model.

15. The computer-readable medium of claim 14 wherein using a language model comprises using a language model that is based on syllable-like units.

16. The computer-readable medium of claim 15 wherein each syllable-like unit comprises a sequence of phonemes and wherein using an acoustic model to score a possible phonetic description comprises generating acoustic model scores for each of the phonemes in a syllable-like unit and summing the acoustic model scores of the phonemes to generate an acoustic model score for the syllable-like unit.

17. The computer-readable medium of claim 12 wherein:

converting the text of the word into a text-based phonetic description further comprises generating a score for the text-based phonetic description based on the correspondence between the text-based phonetic description and the representation of the speech signal;

generating a speech-based phonetic description further comprises generating a score for the speech-based phonetic description based on the correspondence between the speech-based phonetic description and the representation of the speech signal; and

selecting between the text-based phonetic description and the speech-based phonetic description comprises

selecting the phonetic description with the highest score.

18. The computer-readable medium of claim 12 wherein the steps further comprise:

receiving an instruction to generate an audible pronunciation of a phonetic description previously added to the speech recognition lexicon;

retrieving the added phonetic description from the speech recognition lexicon; and

causing an audible pronunciation to be generated based on the retrieved phonetic description.

19. A speech recognition system having a language model generated through a process comprising:

breaking each word in a dictionary into syllable-like units;

for each word, grouping the syllable-like units of the word into n-grams;

counting the total number of n-gram occurrences in the dictionary; and

for each n-gram, counting the number of occurrences of the n-gram in the dictionary and dividing this count by the total number of n-gram occurrences to form a language model probability for the n-gram.

20. The speech recognition system of claim 19 wherein breaking each word into syllable-like units

comprises breaking the words by preferring syllable-like units that occur more frequently in the dictionary over syllable-like units that occur less frequently.

21. The speech recognition system of claim 20 wherein breaking each word further comprises updating the frequencies of the syllable-like units into which the word is broken.